A Pragmatic Approach to Discussing Intelligence in Systems

Gary Berg-Cross
Knowledge Strategies Division, SLAG, Inc.
Potomac, Maryland 20854, USA

Abstract

The nature of the intelligence found in intelligent systems (IS) can be enumerated & framed as categorical dimensions as discussed by as in Messina et al (2002) & Berg-Cross (2002) or built into Agent Ontologies. Berg-Cross's network hierarchy of dimensions include Behavioral/What, Symbolic/Functional Architecture (How), Environmentally Reactive (External why) which is organized by a Goal-oriented and Intentional (Internal why) component. This paper extends the discussion of IS performance in two ways. First, the complex, incremental growth of intelligent functionalities is discussed as a Ptolemaic situation, where multiplies paradigms co-exist but success has been gradual. Relevant examples include commonsense physics and complex problem solving, which are simpler than the full IS problem, but which have afforded only modest progress.

It is argued that in such times there exist fundamental definitional problems which requires more than incremental growth piling up data about functional models of intelligence. Currently several paradigms seem to be under discussion simultaneously. The IS field may be in an hyper paradigmatic situation which requires integration of multiple paradigms more than a shift between paradigms. The second thrust considers the value performance issues of encompassing pragmatic perspective. Aspects of Pragmatic Philosophy's approach to knowledge, truth and reasoning are highlighted. Charles Peirce's pluralistic formulations on reasoning and knowledge are eludidated by Sowa's knowledge soup metaphor and the rational-empirical cycle of reasoning. "Better" measures of IS performance are likely to be found in such pragmatic models that look at "success" in dynamic environment adding cognitive component to traditional KR semantic and syntactic properties . Measurement would then include "Insider" type of Belief, Desire and Intention controlling "social" level metrics, coupled to "outsider" performance metrics.

1. Introduction

Over the last 30 years or so AI and cognitive psychology have been friendly collaborators with distinct but related goals when it comes to discussions of intelligence and intelligence systems. The annual PerMIS workshops provide a useful forum for both AI and Cognitive Science professionals to move beyond isolated phenomena and capabilities to discuss "*complete intelligent systems*". As part of the annual PerMIS conferences several frameworks have been proposed for characterizing the intelligence found in intelligent systems (IS). Two of the broadest, ambitious but practical approaches are found in Messina et al (2001) and Meystel (2001). They enumerate a series of functional "features" ranging from sensing to various knowledge and learning features. Meystel (2001) organizes work around a multi-dimensional knowledge "space" describing intelligent functions. Thus, a system is functionally intelligent because it has "background knowledge" or it is intelligent because it responds appropriately to a stimulus and systems can be located in a vector space of these dimensions. A listing of IS capabilities is illustrative of this approach. Both Messina et al (2001) and Meystel (2001) propose the long list of "properties" that can "tacitly (be) considered to pertain to intelligent systems". Contrasting lists are shown in Table 1 and there is almost no overlap or easy mapping between the items. Messina et al (2001) more closely represents a functional list whereas Meystel (2001) uses a sense-think-act cycle as part of a structured approach to relate knowledge, success and learning. In this model action is not generated directly by perception, but there is a mediated mechanism for behavior generation. A knowledge filled "World Mode" de-couples perception and action and allows "rational" behavior. There are several problems with this approach starting with the use of sequential sense-think-behavior modules. Cognitive studies show perceptual-motor interactions as tasks are learned. Also "knowledge flowing" means very different things (different representation from geometric to feature based to symbolic) in different parts of the "cycle". It is an open issue how knowledge states are represented and whether the meaning of symbols must be grounded in the system's own interactions with the real world.[1] All of which leads to knowledge translation/integration issues

Messina et al 's (2001) list includes:	Meystel et all (2001) faculties include:
•the ability to deal with general and abstract information •the ability to deduce particular cases from the general ones •the ability to deal with incomplete information and assume the lacking components •the ability to construct autonomously the alternative of decisions •the ability to compare these alternatives and choose the best one •the ability to adjust plans in updated situation •the ability to reschedule and re-plan in updated situation •the ability to choose the set of sensors •the ability to recognize the unexpected as well as the previously unknown phenomena •the ability to cluster, classify and categorize the acquired information •the ability to update, extrapolate etc.	•Allows an agent to deal with knowledge to achieve externally measures success externally under a particular goal . "Knowledge" of an AGENT is the collection and organization of information units. "presumed to appear as a result of *Learning* about the objects of the external world, their (objects) interconnection, and processes of changes produced by the AGENT within this external world . Learning process is understood as recording the *experiences* encountered by an intelligent system and deriving from these experiences a new set of rules that suggests how the intelligent system should act under particular circumstances . Experiences are understood and stored as triplets of the information units "situation ->action ->newsituation" that allow the behavior generation module of the AGENT to infer what is the action that is required to improve the situation

Table 1 Contrasting IS Functional Lists

Structuring functions into a coherent model by means of theory is one way of improving over the simple lists. Meystel (2001) proposed a general intelligence design, depicted in Figure 1 (after Gudwin 2000), consisting of six "consecutive functional elements connected by a flow of "knowledge".

[1] Still another mix is classic symbol processing AI functionalities with non-symbolic connectionist approach.

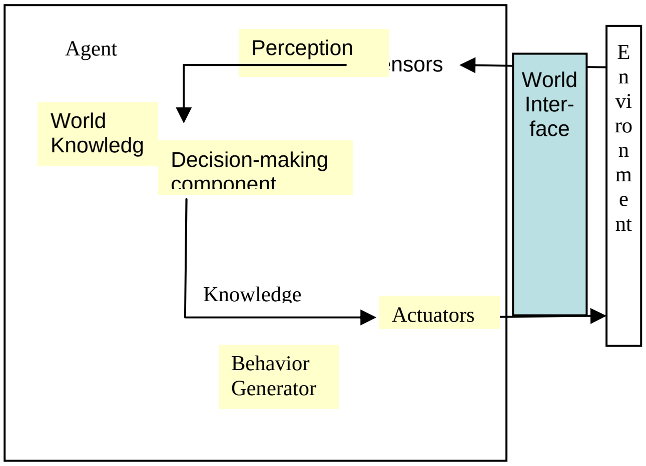

Figure 1 After Meystel (2001) Model of an IS

This diagram follows a traditional theory to explain the behavior of an intelligent system:

> "Intelligence is a faculty of the system that provides an ability of a system to act appropriately in an uncertain environment, where appropriate action is that which increases the probability of success, and success is the achievement of behavioral sub-goals that support the system's ultimate goal."
> Albus (1991)

It is intriguing to note how this model differs from Shannon's communication theory. That theory-based model consists of six elements: information source, transmitter, channel, noise source, receiver and destination which are selectively linked by directed relations according to the direction of communication. In the theory communication takes place when a message produced by the information source is conveyed into the channel by the transmitter. In the channel the message may be distorted by signals from the noise source. At the other end of the communication system the receiver receives the signals from the channel, passing them on to the destination for final processing. With this description Shannon conceptually summarized those aspects of information falling under the heading of functional-cybernetic information. Moreover, the accompanying mathematical-statistical formulation of the processes forms the formal basis for a technical realization of communication systems and computers. Yet Shannon deliberately excluded important aspects of information such as semantics because they seemed irrelevant to the technical realization of his concepts.

Contrast this with Albus' (1991) model. As noted by Gudwin (2000) this is very similar to a theory proposed by the Philosopher Charles Peirce's and is a pragmatic definition. It presupposes an agent governed by goal-oriented *decisions* based on *knowledge* and it presupposes the system is able to organize acts in a way to achieve this goal. As such it employs some semantic principles to rise above Shannon's information theory (IT). IT owes some of its success to the fact its formulas can immediately be applied to practical problems, but I believe that the omission of a semantic aspects limits it utility. Albus' (1991) construct builds on the ubiquitous and parsimonious cognitive idea of a "rational agent' using 'knowledge' to 'succeed' in the world. Various studies and traditions have qualified the notion of rational agents making them more bounded, approximate, and non-monotonic. Taken as a whole, it is clear that rationality is neither omniscience or optimal. More accurately agent rationality is always limited by processing, perception and knowledge and thus we might speak of it as pragmatically seeking to optimize expected**,** not actual utility. One of the conflicts in the IS field has been how central these aspects are to a cognitive agent. Gudwin (2000) notes that a difference between Albus & Peirce is that the Albus' definition doesn't impose the necessity of having a measure of how much of a goal is being achieved, a topic related to pragmatic definitions of success.

All of which serves to illustrate the distance the IS field is from reaching convergence. The situation is one of accumulated complexity. We can see part of the problem when we look at ontological definitions of some of the key typically cited concepts: knowledge, success, appropriate rational, task of defining such words raises questions that involve almost every other aspect of ontology (Sowa, 2000).[2]
Consider the concept of rational agents as those agents that perform rational actions. Which are in turn requires the meaning of doing the "right" action. The "right" action is one that will cause an agent to be the most successful. "Success"? Can we define it as a simple *Performance measure?* Well there are several simple performance measures for a cleaning robot (see Gunderson and Gunderson, 2003 for a discussion of related issues on the performance of a cleaning robot). There are immediately pragmatic issues of what we mean by successful cleaning. Is it the amount of dirt cleaned up by robot? Is this relative to total amount and standards?
Is it the number of clean tiles on a floor or the number of clean tiles after each circuit? Which is better? Pragmatic philosophers point out that we can design success according to what is wanted or how one thinks an agent should behave. If we design it as we want it we are not giving the agent the possibility of rational awareness of how goal performance is being measured. We might agree with Pragmatists like Charles Pierce that this is a missing aspect of intelligent behavior.
Peirce's ideas on the components of a rational process this is taken up more fully in Section 3.
One popular assumption affecting the field's strategy is the idea that intelligence arises from the interaction of a very large

[2] An even stronger argument over linguistic problems with IS concepts is made by Cottam et al. (2003) in this Proceedings.

number of basic components and thus there are few simultaneously general and powerful laws for intelligence. The implication is that we need to cast a large net of many functions and expect many "micro" principles to apply. However to philosophers like Fodor (2000) an enumerative functionality strategy is a sign of flaws in the "normative" conceptual tools of IS. The underlying flaw is that making functional mappings of inputs to outputs part of a symbolic processing paradigm that using many implicit assumptions about central problems, provides no real explanation. To take an overly simple example, consider the functional explanation of intelligent systems' ability to "cluster, classify and categorize the acquired information". One may postulate the existence of a "set of classifying functions" which selects appropriate input. Such explanations are arbitrary starting points if not "vacuous. It is a natural human tendency to ascribe mental faculties to a system in such a direct way to bypass legitimate philosophical issues. Going the other direction away from functional approaches to pure subsumption architectures doesn't fully handle this problem. As Gudwin (2000) argues , while subsumption practically situates an agent and handles the behavioral aspects of intelligence, there are severe constraints in such applications. For one thing such systems don't explicitly know its objectives and thus cannot provide a guaranty that goals will be met. They also depend on an adequate specification of the set of available behaviors.

One way of to gauge IS progress is to look at extensive work in related areas that are both simpler and have started earlier. Naive Physics (Hayes, 1985) and Complex Problem Solving (Funke 1995) are useful examples. It is almost twenty years since Pat Hayes' "Naive Physics Manifesto" first proposed a research program to formalize and implement physical reasoning at the commonsense level –a useful "module" for a complete IS. Generally, the proposed methodology for representing commonsense knowledge used first-order logic & its variants, while assuming the existence of an automatic deduction system which would provide "heuristic adequacy" given the epistemological strength of the representation. This type of commonsense reasoning would be an essential part of a practical IS giving it a capability to understand, and interact with, the everyday world around it. As a strategy this is a much simpler goal than building a full IS including agent abilities where cognition and intentions adds considerable complexity. The experience in Naïve Physics, however, has been one of very slow progress and as noted by Davis (1999) beset with many foundational issues. For example, how do we define & scope naïve physics? How do we handle intuitions or agent intentions? Over time researchers have tried sophisticated representations beyond first order and modal logic, but the problems of common sense have not yielded easily. A review by Watt (1997) found several flaws in the work including: ignoring control issues (which hasn't proved successful in practice); the assumption that representations can be separated from the use of those representations (which is also suspect), and the need for some measure of heuristic adequacy. In critiquing Hayes' strategy of 'micro-worlds' (small isolated domains such as the behavior of different forms of liquid) to make common-sense reasoning tractable, Watt (1997), notes that this seems to have merely pushed many of the real problems into a future integration domain integration task –something that few have tried.

Complex Problem Solving (Funke 1995) is another related effort to model thinking and reasoning that started with experimental tasks (also called Microworlds as in Niave Physics) that are dynamic, time-dependent, and complex. Firefighting and soccer are examples of such problem domains. Compared to traditional "Problem Solving", Complex Problem Solving (CPS) radically changed the kind of phenomena reported (in the real world states are continuous with no uniquely defined states); the kind of explanations looked for, & even the kind of data that is generated. Information is uncertain/ incomplete, dynamic and unpredictable. As noted by Quesada et al (2002) the results obtained to date in CPS are:

> " far from being integrated and consolidated. This fact led Funke (1995) to affirm that 'Despite 10 years of research in the area, there is neither a clearly formulated specific theory nor is there an agreement on how to proceed with respect to the research philosophy. Even worse, no stable phenomena have been observed' ... Almost another 10 years after Funke's argument, although more empirical research has been conducted in the area, we cannot say that the situation has changed drastically. At this moment, there is no theory able to explain even part of the specific effects that have been described or how they can be generalized. A theory of generalization and similarity is as necessary to psychology as Newton's laws are to physics."

Taken as a whole I speculate that the current state of work (in both IS and related, if simpler, problem fields) seems a bit like a Ptolemaic paradigm. That is, aggregate functional models are a bit like Ptolemaic approaches as they break complex intelligent behavior into neat, related components that are analogous to "perfect circular motions". It is a system of cycles and epicycles with certain perfect cognitive process as the functional/circular primitives that we are willing to accept as heuristic devices for their practical computational usefulness. Such core functions are added all to in an ad hoc way as required to obtain any desired degree of performance and accuracy. Using large combinations of constructions we are able to measure performance in some small problem domain for some of the agreed upon intelligent behaviors within the standards of observational accuracy. The concern is that, like a Ptolemaic system, the model is more an after-the-event description then a deep capture of the underlying problem. Such descriptions may be very useful, making it possible to economically summarize a great amount of brute observational data, and to produce empirical predictions, but they often prove to be brittle, do not scale and might not be fruitful to further predictive IS research. To illustrate further, consider cognitive research functional "components" as in Figure 1. They aren't really independent – they are highly interdependent. Systems built with such modules have many questions that can not be answered as well answers that are extremely complicated due to component interactions. For example Schoelles & Gray (2000) studied interactions between perception, cognition and motor performance on target tasks. They used the ACT-R/PM architecture (Anderson & Lebiére, 1998) which combines ACT-R's general theory of cognition with modal theories of visual attention (Anderson et al 1997) and motor movement (Meyer & Kieras 1997). In time critical situations selecting "threat value objects" Schoelles & Gray (2000) needed to propose additional/ad hoc microstategies (epicycles?) by which subjects manipulate interactions between perception and motor acts to improve performance. Similar qualifications and epicyclical additions have been advanced about the structure of agent knowledge and useful meanings of "success".

Using a Ptolemaic metaphor affords a discussion of how particular models arise from more general theory. Classically one paradigm has largely ignored other paradigms talking past issues along with anomalies in observations they predict. There seems to be plenty of room in alternate approaches to IS for such ignoring of competing theories. Thus in the current situation, alternate paradigms aren't easy to knock down and reject. They hold on. Perhaps in IS we need an integrating hyper paradigm to provide points of contention and agreement.
Switching to multi-paradigm integration might involve a new "world view". It should be rational and pragmatic.

The remainder of this paper is organized as follows. First basic pragmatic principles, positions and issues are reviewed as context. Then pragmatic definitions of knowledge are elaborated and related to Sowa's Knowledge Soup concept as a contemporary inheritor of the pragmatic enterprise. A final section concludes with some implications and directions.

2. The Pragmatic View of Success, Instrumentalism and Intelligence

The Pragmatism philosophical movement was developed over a century ago as an American approach to the great problems of philosophy such as the truth & knowledge. Pragmatism proceeds in a modern, logical but practical way and has tended to criticize traditional philosophical outlooks (Descartes, Kant etc.) in the light of scientific and social developments. To pragmatists most propositions and questions discussed on philosophical matters were not so much false as senseless and lacking of meaning. Pragmatism's basic position holds that both the meaning and the truth of an idea is a function of implications that are often described as practical outcomes. Thus, in discussing good and evil pragmatists saw it as dependent on its practical effects (or success) on human behavior.[3] Fundamental to pragmatism is a strong anti-absolutism tenet helpful to Science: the conviction that all principles are to be regarded as working hypotheses rather than as metaphysically binding axioms. Charles Sanders Peirce (All citations to Peirce are from Hartshorne, C. and Weiss, P. Collected Papers of Charles Sanders Peirce, Volume 5 and Volume 6, Harvard University Press, 1965) is considered the founder of pragmatism[4] and his views of truth, reasoning and knowing provides a foundation for discussion. Peirce developed a theory of meaning in the 1870s, holding that there is an intrinsic connection exists between meaning and action -- that the meaning of an idea is to be found in its "conceivable sensible effects". Peirce might say that we can understand a concept's meaning by looking at the real world implications of various hypotheses about it. But this is far from a simple utility view as shown by Peirce's elaboration:

> "the true meaning of any product of the intellect lies in whatever unitary determination it would impart to practical conduct under any and every conceivable circumstance, supposing such conduct to be guided by reflexion carried to an ultimate limit." Peirce (1902)

Thus meaning is a belief hypothesis that is provisional, dynamic and fallible. Peirce centered thinking in human-generated belief which arises through what he called "habits of action." Peirce would agree with a direct implication of Figure 1 - that the core function of thinking is as a stage in the production of successful action in an environment. And like Berg-Cross's (2002) 3-level model, a key feature of Peirce's view of intelligent behavior is an essential connection between "rational cognition and rational purpose"[5]. In proposing an essential relation between intelligent thought and human conduct, Peirce was not subordinating reason to action, utility/profit or individual interests. Rather, Peirce defines a concept 's meaning as a form which is most directly applicable to self-control in any situation and to any purpose. To him, the rational meaning of a proposition is its future potential which he regards as the ultimate test of what truth means. If this definition seems vague Pierce indeed found the

---

[3] The argument is further that we see in nature only that narrow range of physical parameters which are compatible with our evolution as complex, self-reproducing, learning organisms.

[4] In conflict with James' popularization and simplification of the idea, Peirce later changed the name of his philosophical position to "pragmaticism".

[5] A striking feature of Peirce's theory was its supposition of "an inseparable connection between rational cognition and rational purpose; and that consideration was which determined the preference for the name *pragmatism*"

concepts of “real” and “true” to be fundamentally slippery. [6] Peirce’s position on these concepts is well reflected in John Sowa’s (2000) Knowledge Soup metaphor, which is discussed in Section 3. In this view truth is not so much fuzzy as it is context dependent. It is a relation between a theory and a model of some aspect of the world for some purpose. For example, we may ask is a pizza circular or is a ball bearing circular? The answer is yes or no, depending on our purposes.

The development of Instrumentalism by John Dewey may be seen as an elaboration on this observation. Dewey, like Peirce criticized the traditional notions of truth embodied as falsely precise logical theory of concepts, judgments and inferences in their various forms. Thus Dewey made the inquiry process, rather than truth or knowledge, the essence of logic and of Science. Dewey, more than Peirce, took the time to explore the larger social implications of a pragmatic view of truth and thus Instrumentalism prefigures the Kuhnian social view of scientific truth in several ways. Given the prior reference to Ptolemaic models, a digression to Kuhn is illustrative since he uses the Copernican revolution from Ptolemy’s thought as his driving example.

Kuhn, like Dewey, focuses on process and finds a basic pattern, called a paradigm, from which scientists see and interpret the data within a particular field of empirical enquiry. Finding a new pattern/paradigm, as Copernicus did, can cause sudden, irreversible shifts in understanding. This reflects the social nature of scientific truth. Thus, once we have seen a good, new paradigm, and the relevant data in light of it, it is not easy to ignore it, and to return to the state of seeing just through the old paradigm because paradigms are, at least initially incommensurable. Which is to say, that competing theories involve different schemes of organization that allow no common standard of measurement. Taking IS performance as the topic then, one may hypothesize that until we agree on the outline of the whole of “Intelligence” about whose components we are arguing, we have inadequate common ground for communication, since Kuhn argues that 'parts' are relative to their wholes and hence have no meaning (or a different one) if taken out of their context. I attribute some of the difficulty in the areas of Naïve Physics and CPS to the absence of goals & intentions in the meaning of core concepts. We have similar challenges in defining IS performance.

In concluding this section we note that the broad influence of pragmatic approaches of the early 20th century waned as people like Church and Turing developed classes of computable functions and viewed these operations as analogous to what a mathematician does while "computing" a function in the sense of evaluating it by application of a rote procedure. This powerful declarative approach blossomed into a computational theory of intelligence and seemed to leave Pragmatism behind as it moved on to correspondence & coherence theories of truth. However, pragmatic discussion seems to be on the increase as difficult problems are addressed. One may find pragmatic approaches in collections like Giere (1992) who notes that "the only form of rationality that exists is the instrumental use of empirically sanctioned strategies to achieve recognized goals". Pragmatic principles are also seen in arguments by Fodor (2000) ranging from of the need for pluralism to problems in the semantics of mental States. Indeed the spirit of PerMIS work is solution-driven, usually starting from some computational mechanism (such as symbol processing or neural nets) and arguing what phenomena it can account for. I believe there are several reasons that such pragmatic formulations are back, including the discomfort of an increasing Ptolemaic situation along with the realization that foundational problems that beset cognitive science itself. Also a degree of pluralism and relativism is more fashionable in part because past criticisms, such as Wittgenstein’s, have never been addressed. For example, Wittgenstein argued that most philosophical questions and propositions, such as underlie our scientific theory, result from the fact that we do not understand the logic of our language, which we use to describe, formulate and pass on in Science. If this view is correct the IS field is simultaneously using several practice paradigms which makes it is difficult, if not impossible, to agree on key issues. Kuhn argued that by naively trying to win an argument one first acknowledges some aspect of the paradigm one is wishing to leave behind. But the plausibility, for example, of Copernican cosmology could never be demonstrated in terms of the Ptolemaic system. Typically what is needed is for the new paradigm first to be adopted in the hope that it will commend itself as a more adequate or satisfying frame within which to make sense of the relevant data. I take the implications and guidance for this to mean that the incremental, feature adding approach to the complexity of an IS is now faced with several new concepts and terms about intelligence, not addressed in and competing with the modal model of IS – the original symbol processing approaches. This paradigm is in competition with a Peircian view of knowledge, the Connectionist approach, evolutionary systems and new fuzzy additions to cognition and language that Zadeh (2002) summarized at PerMIS. Connectionism, for example, embodies a very distinctive characteristic to distinguish cognitivism without repudiating cognitivism itself. Instead it is simply providing an alternative to the standard rules and representation view of cognition. But it is an open question whether it an adequate theory without proposing some symbolic processing in addition to sub-symbolic processes Zadeh (2002, 2003) takes a completely different tack, espoused the centrality of abstract, generalized knowledge that in turn affords new IS formulations such as

[6] John Sowa quotes Peirce’s conjecture that “.. truth is rather on the side of the Scholastic realists that the unsettled is the primal state, and that definiteness and determinateness, the two poles of settledness, are, in the large, approximations, developmentally, epistemologically, and metaphysically. “

perception-based language for computation He argues, for example, that it is not possible to formulate a general definition of causality within the conceptual structure of traditional, bivalent logic. Instead causality is an inherently fuzzy concept because it is always a matter of degree. Likewise he proposes a "Precisiated Natural Language" based on a perception-based theory which gives an agent the capability to operate on perception-based information. Generalizations from this theory are significantly more complex than traditional "measurement-based versions". In this instance, as in many others, Zadeh argues that an up front complexity is the price that needs to be paid to reduce the gap between theory and reality.

With this as background the next section turns to continuing problem with the nature of knowledge.

And describes John Sowa's reformulation of Peirce's view on reasoning and knowledge.

3. Knowledge Soup & Pragmatic Reasoning

Knowledge plays a key role in traditional computational view of cognitions as well as new paradigms as described above. Traditionally knowledge representations have both semantic and syntactic properties, but processes of "reasoning" are responsive only to the syntax of the symbols. In this view "computation" is formal symbol manipulation (I.e., manipulation of symbols according to purely formal--i.e., non-semantic--techniques.) and is equivalent to reasoning. The computational account of cognition depends essentially upon a prior claim that intentional "states" involve symbolic representations. Accordingly, these representations have both semantic and syntactic properties, and processes of reasoning are performed in ways responsive only to the syntax of the symbols--a type of process that meets a technical definition of 'computation', and is known as formal symbol manipulation. (I.e. manipulation of symbols according to purely formal--i.e., non-semantic--techniques. Peirce challenged just such simple concepts of context free propositional meaning, seeing meaning as a flexible, contextual form most directly applicable to self-control in any situation and related to intention and purpose.

Newell (1982) moved the AI field in this direction by introducing a separate, higher level knowledge concept to the symbol-processing formulation making it distinct from the symbol level in which representation (e.g. logic) lies. Newell also emphasizes an abstract but pragmatic view in proposing that this view of knowledge permits prediction and understanding behavior "without having an operational model of the processing that is actually being done by the agent." This is a richer model of knowledge still slowly maturing, but it is worth looking at such knowledge in actual systems as exemplified by Knowledge-Based Systems (KBSs). Do they measure up to these concepts of goal-oriented knowledge? One way to explore that question to consider user experiences with KBSs, their reasoning and knowledge. At a first glance we see that the early formulations of knowledge bases presupposed an answer to an old philosophical question. That is, the typical knowledge engineering approach assumed that knowledge is available as some type of frozen thing existing in experts and transportable unchanged to the knowledge engineer for storage in a knowledge representation. One sign that such a knowledge chain (from expert mind to internal storage) has problems is that the resulting knowledge often appears incomprehensible to end-users. Users site differences in meaning based on context and sometimes use fuzzy concepts as previously noted. This is one of the reasons that KMSs are often rejected (Woods et al., 1990). Brezillon and Pomerol (1996) discuss the KBS rejection problem from the point of view of the social nature of an "intelligent" medical diagnosis. They argue that when the expected output of a system is a diagnosis, healthcare workers orient to this output as a confirmation of or as a disagreement with their own diagnosis "hypothesis". If there is a conflict they expect to enter into a dialog to solve the conflict. But KBSs don't typically include the functionality of explanation, understanding dialog etc.[7] This is vastly beyond the network of their isolated expertise, which may answer "What is the fault?" or "What is the problem?" Kidd and Sharpe (1988) also note that users of diagnostic KBSs more typically are seeking answers to formulations such as "Why did fault A happen?", "Will remedy B fix Problem Y?" or "Can I test C without affecting the level of D?" A Peircian view might offer 2 suggestions. First as to the knowledge in the system, he would recognize that captured, propositional knowledge lacks connections to uses of that knowledge to achieve goals (such as described in Berg-Cross, 2002). Second at the level of system, Peirce might argue that the behavior of such KBSs is not really intelligent, because they cannot operate on their own knowledge to refine the meaning through social (system and human) interaction over time.[8]

A still stronger critique of the knowledge in ISs/KBSs concerns the knowledge in World Models, such as depicted in Figures 1 or 2. World models are useful for many things including as affordances to navigate in the physical world and handle objects.   World models should also allow an IS to anticipate the results of their own actions, a process simpler for interacting with non-agent objects than with Agents. Pomerol (1995) notes that many typical industry process

[7] Zadeh might argue that one reason is that perceptions play a key role in human cognition but not in current machines implementations of intelligence.

[8] Dewey anticipates some K-Rep issues as a "verification" issue. Verification of knowledge is not as a passive "looking-at" the sensibilia "given" in experience --rejecting this as another example of the hold of "specator theories of knowledge". Rather, verification is an integral part of the process in which human agents interact and cope with problems that are thrown up by their environment --practical, rather than theoretical problems.

KBSs assume a simple, direct model relating action to results. This may be appropriate for such processes where there is a controlled environment and the average time between actions narrows the degree of uncertain change that might develop.
In Berg-Cross (2003) a more elaborate world model gets coordinated by agent plans and belief modeling using a Belief-Desire-Intention (BDI) agent model. A BDI model structures knowledge within an agent – in terms of beliefs, goals and plans – which simplifies communication and allows coordination between agents. Zadeh (2002) would add that the fuzziness of causality is not addressed by traditional implementations and thus are applicable to only a subset of situations. Outside of these realms increasingly fuzzy models of causality may be needed. A uniquely Peircian view of world model knowledge and its relation to pragmatic reasoning has been constructed by Sowa (2000) to which we now turn.
The general pragmatic framework allows for several kinds of reasoning and with supports for declarative knowledge, explicit definitions, procedural knowledge and the idea of family resemblance. Sowa (2000) organizes Peirce's Rational-empirical reasoning into three parts:

1. Induction or learning which start with observations and looks for commonalities (a basic cognitive process) deriving a theory to summarize the data.
2. Deduction or inference which starts with a theory and observes some new data. Then uses the theory to logically generate implications
3. Abduction or guessing which starts with disconnected observations and guesses (hypothesizes) a theory that relates them. The test of this hypothesis is by means of pragmatic tests using subsequent induction and deduction.[9]

These processes are organized into an overall system of (Figure 2) by Sowa using an additional idea - knowledge soup. Knowledge soup captures the idea that real human knowledge is fluid and lumpy, with adherable chunks of theories and hypotheses that float in and out of awareness. As knowledge circulates through this process it becomes more meaningful, truthful, validated and thus better corresponds to models of reality and is more coherent.
The process allows for a natural history of knowledge flowing from conjectures and theory to prediction and groundings in observations both at the formative and subsequent part. Taken together it structures Peirces idea true meaning of any "product of the intellect" as that which would impart to practical results under "any and every conceivable circumstance, supposing such conduct to be guided

[9] An implication of such a model of reasoning is that science comes naturally to us.

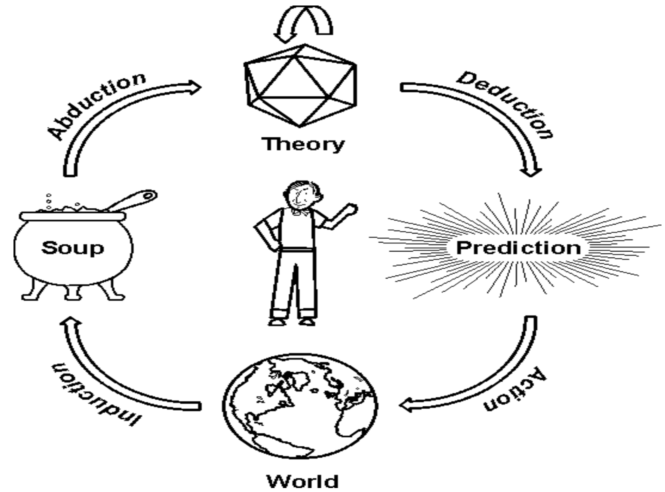

Figure 2 Sowa's summary of Peirce's rational-empirical reasoning theory

by reflexion carried to an ultimate limit." It is worth noting that there are variety of natural and social constancies that make meaning converge to allow correspondence to both the world and between individuals.
Figure 3 expands Sowa's discussion to further `illustrate relations between knowledge soup contents and the world. Reality stands on the left while on the far right is a lumpy part of our knowledge soup externally expressed as formal theory in represented in predicate calculus such as Peirce pioneered. But as Johnson-Laird (2001) observed, brains do not use Logic, they use "mental models". In the middle is more recent formulation of such a larger Tarski-

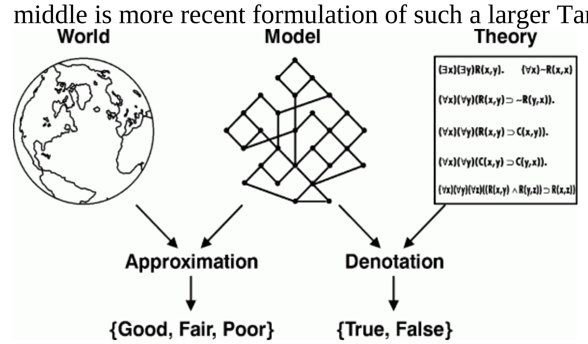

Figure 3 Denotation and Approximation Mappings between Models, Theory and Reality (Sowa 2000)

style model expressed by Sowa as a graph. World models are also part of our knowledge soup but their expression is more abstract and approximate. In the model nodes represent individual model entities while arcs represent relationships. The figure represents a part of what we might call pragmatic model knowing represented as a mapping from the symbols of the formal theory to vaguer symbols of the model. Sowa proposes that this mapping determines the denotation of propositions as having a binary {true, false} truth value. We might speak of these as our formal abductive hypotheses about the larger model. Even the symbols and logic chosen to represent these may be considered hypotheses of what will best represent the larger, soupy model.[10] The semiotic nature of the symbol to symbol mapping is an inherent feature of this view. However the mapping from the model symbols to certain aspects of the world relates to the pragmatic adequacy of the model as an approximation designed and intended for some particular purpose. But agent purposes are endless and reflect the situatedness of our cognition so there is no simple, ultimate model. As a consequence when we express the central Tarskian model in formal mathematical structure we are forced to use some abstract simplifications of tokenized symbol types. These are abductive hypotheses at a more abstract level. Language is not shown in the diagram, but we see it as a more informal attempt to express a portion of the model in the way a spoon might take a portion of the soup, lumps and fluid together. This seems to be a useful, social and intelligent adaptation of limited, rational intelligent to the complexity of the real world. Language is obviously symbolic, and we share these symbols over groups. Language has pragmatic meaning to the extent that some conventions hold within the linguistic community, but different people have different knowledge soup meaning background or , views, theories, purposes, and vocabulary. For Peirce every natural language utterance has three components, (*s*,*p*,*b*): Speech act *s*, which states the purpose at a performative/intention level, a Proposition *p*, which states the content at what we might call a conscious level and Background knowledge *b*, which is unstated, but assumed as the fluid within which the nuggets of propositions fit. Thus a speaker may imply many things more than in the propositions of a sentence. These are its full "meaning" including the purpose for which a sentence is "used" (Grice, 1975).

There are a number of interesting relations between Figures 2 and 3 worth making briefly in this regard, not only for individual agents but for communities of agents who share theories and models through formalisms and language. Knowledge soup gets populated with connected pieces as an agent's hypotheses and theories flow around the empirical reasoning cycle. Validated theory and models get expressed and shared. Thus Albus and Meystel's (2001) reference model architecture expressed in terms of the Real-time Control System (RCS) may not be all the knowledge soup there is to the IS issue, but it solidifies a significant portion of current theory and practice to achieve a coherent model for the design of intelligent systems. This is useful as tests of approximation to reality. It is also worth noting in passing that the rational-empirical reasoning theory provides a natural framework such things as Piagetian assimilation and accommodation such as addressed by Arata (2003) in this Workshop.

It is also worth noting that Sowa points out several different places where vagueness could be located in the model. His chief argument is that there is a " fallacy of misplaced fuzziness", when fuzzy theorists associates fuzzy numbers with the truth values determined by the mapping to theories on the right. Sowa's formulation sees fuzzy numbers as better applied to the approximation on the left – mapping from our model to reality. It is in this language that we have the vagaries that Peirce alludes to.

Among other things the knowledge soup and rational-empirical reasoning theory helps us understand certain problems in acquiring knowledge for "bases". We are all too often capturing nuggets with the soupy fluid. As a by-products of our mental models we make mistakes because of the efficient way but ineffective way we represent and reason about the world. There is an analogous knowledge problem in our ISes. Problems with simple knowledge bases have lead to more sophisticated efforts to capture knowledge within a larger context that provides its relevance fluidly across related situations. For example (Russell & Norvig, 1995, pp.20-21) discuss the challenge of capturing "background knowledge" functions. In general IS implementations lack background knowledge and/or run into intractable computations to make up for limitations in knowledge representation structures. There is an extensive literature on such addressing such problems as situated knowledge and the role of context in knowledge. This literature often combines recent cognitive theory with a pragmatic philosophy framework that includes the environment as a major factor in intelligent systems. In such efforts, functions need to be supported in a wide context including inter-agent environments that are developmental drivers of intelligence.

4. Conclusion

This paper has extended the discussion of ISes by exploring the idea of pluralistic intelligence in a fashion based on pragmatic philosophy's approach to the nature of intelligence and truth. Two types of implications are worth noting, one at the individual, system level and the other at the level of Science, its goals and practice. At the system level internal, pragamtic realists like James and Dewey believed that intelligence, like the external world, can be"correctly/ pragmatically" described from a number of different perspectives. These views and hence the "meaning" of things reflect an agent's internal interests and purposes. Thus external definitions at the behavioral level are not adequate in isolation. An essential part of the pragmatic view have been implemented as part of BDI architected agent systems. BDI architectures provide agents goal-directed behavior. And this notion of intentions is practical because it implies some

[10] Sowa discuses written symbols as representing axioms of a theory, whose implications are some abstract set of propositions. Peirce might note the pragmatic semiotic nature of such mappings as part of the task to ascertain laws by which intelligence uses one sign to generate another, as "one thought brings forth another." Peirce saw pragmatic semiotics as the study that relates signs to agents who use them to refer to things in the world and to communicate their intentions about those things to other agents who may have similar or different intentions concerning the same or different things.

commitment to these actions to achieve a goal. Further. As Norling et al (2000) argues there is a very pragmatic basis for intentions – they help prevent complicated reasoning at every time step, since once an agent has decided to do something, it will continue to do it until it becomes either impossible or unnecessary. This applies both to physical actions and to plans to test hypotheses about the world as Piaget might have described such cognition.

Perhaps Peirce had the clearest things to say about the balancing of the goals of science and the fallible role of belief in its practice. We have a goal of describing intelligence but not even the most advanced scientific one, is Nature's final one. As is true of other scientific attempts there is no simple path to the truth about intelligence. Descriptions of intelligence as available to us, even as scientists, are grounded in our purposes & paradigms as professional investigors. Our practices, their truthfulness and reference are actually "internal to conceptual schemes serving different purposes." This position is well captured in Sowa's knowledge soup metaphor and needs further illustration to determine its usefulness.

Taken together one direction of the argument leads to the idea that our approaches to IS work seems a pre-scientific Ptolemaic one, in that there is no consensus on theory and we have many incompatible/incomplete theories. For example, there are fundamental issues of how agent knowledge relates to the world. As a result differing approaches are challenged to agree on key experimental performance measures of mediated behavior. We may be able to use some traditional, if somewhat lose, criteria of scientific explanation to help, including symmetry, elegance, and simplicity to select performance measures. I see these as abstract abductive hypotheses such as discussed in the prior section. For some researchers fuzzy and connectionist approaches seem to provide a firmer base for knowledge, but without a connection to symbols these seem incomplete. At the level of Science it is quite possible that we are in a new Ptolemaic world where there is a fundamental mismatch of the various connectionist, fuzzy & traditional symbol-processing systems concepts and we need a new scientific synthesis. If so, some substantial effort may be needed to find an interaction language between these realms. We may need an integrated, hybrid paradigm such as bridges the quantum and mechanical level in physics. Not an easy task, but an intellectually interesting one and Sowa's knowledge soup synthesis may represent a useful framework for some integration.

References

1. Albus, James S., "Outline for a theory of intelligence", *IEEE Transactions on Systems, Man, and Cybernetics* 21(3):473-509, 1991.
2. Albus James & Meystel Alexander M. Engineering of Mind: An Introduction to the Science of Intelligent Systems Wiley, 2001
3. Anderson, J. R., & Lebiére, C. (Eds.). (1998). *Atomic components of thought*. Hillsdale, NJ: Erlbaum.
4. Anderson, J. R., Matessa, M., & Lebiére, C. ACT-R: A theory of higher-level cognition and its relation to visual attention. *Human-Computer Interaction, 12*(4), 439-462. (1997).
5. Arata, Luis. Interactive Measures and Innovation, PerMIS 2003. September 16-18, 2003
6. Berg-Cross, Gary Dimensions of Intelligent Systems, PerMIS 2002., 2002
7. Berry, D.C., & Dienes, Z. The relationship between implicit memory and implicit learning. *British Journal of Psychology, 82*, 359-373. 1991
8. Brezillon P. and Pomerol J. "Misuse and Nouse of Knowledge-Bases Systems :the past experiences revisited", in Implementing Systems for Supporting Management Decisions, 1996
9. Cottam, Ron, Ranson, Willy, Vounckx Roger. Abstract or Die: Life, Artificial Life and (v)organisms PerMIS 2003. Sept 16-18, 2003.
10. Davis Ernest The Naive Physics Perplex, AI Magazine, Winter, 1999.
11. Fodor, Jerry. 2000. *The Mind Doesn't Work That Way.* MIT Press.
12. Funke, J.: Complex Problem Solving: The European Perspective. Hillsdale, NJ:Lawrence Erlbaum, 1995.
13. Grice. H. Paul Logic and conversation. In P. Cole and J. L. Morgan, editors, New York, 1975.
14. Giere, R.N. (1992) (Ed.) Cognitive Models of Science. Minnesota Studies in the Philosophy of Science, volume 15. Minneapolis: University of Minnesota Press.
15. Gudwin R.R. "Evaluating Intelligence: A Computational Semiotics Perspective" - 2000 IEEE International Conference on Systems, Man and Cybernetics - SMC2000 - Nashville, USA, 8-11 October, 2000, pp. 2080-2085.
16. Gunderson, J., Gunderson, L., Mom! The Vacuum Cleaner is Chasing the Dog Again!, PerMIS 2003. September 16-18, 2003.
17. Hayes, P. "*The Second Naive Physics Manifesto,*" in J.R. Hobbs and R.C. Moore (Eds.) Formal Theories of the Common Sense World (Vol. 1). Norwood, NJ: Ablex Publishing Company, 1985.
18. Hartshorne, C. and Weiss, P. Collected Papers of Charles Sanders Peirce, Volume 5 and Volume 6, Harvard University Press, 1965
19. Johnson-Laird, P.N. Mental models and deduction. *Trends in Cognitive Science,* 5, 434-442. 2001
20. Messina,E., Meystel, A. and Reeker, L. Measuring Performance and Intelligence of Intelligent Systems, paper at PerMIS 2001.
21. Meyer, D. E., & Kieras, D. E.. A computational theory of executive cognitive processes and multiple-

task performance: Part 1. Basic Mechanisms. *Psychological Review, 104,* 3–65. (1997)

22. Meystel, A., Theoretical Fundamentals of Performance Evaluation in Intelligent Systems, PerMIS 2001
23. Newell, A. The Knowledge Level. Artificial Intelligence, 18 (1), 1982.
24. Norling, Emma; Sonenberg , Liz; and Ralph Ronnquist Enhancing Multi-Agent Based Simulation with Human-Like Decision Making Strategies, In Scott Moss & Paul Davidsson, (Eds), *Multi-Agent-Based Simulation, 2nd International Workshop, 2000 Boston.*
25. Peirce, Charles 'Neglected Argument for the Reality of God', 6.490, 1908.
26. Quesada, J., Kintsch, W., and Gomez E. A theory of Complex Problem Solving using Latent Semantic Analysis. In W. D. Gray & C. D. Schunn (Eds.) Proceedings of the 24th Annual Conference of the Cognitive Science Society Fairfax, VA. Lawrence Erlbaum Associates, Mahwah, NJ. pp. 750-755, 2002.
27. Russell, Stuart, and Norvig, Peter, *Artificial Intelligence: A modern approach*, New Jersey: Prentice Hall, 1995.
28. Schoelles Michael J & Gray W. "Modeling Embodied Cognition in a Complex Real-Time Task, COGSI 2000.
29. Schoelles, M. J., & Gray, W. D. Argus Prime: Modeling emergent microstrategies in a complex simulated task environment. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modeling* (pp. 260-270). Universal Press 2000.
30. Sowa, John F. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks Cole Publishing Co. 2000.
31. Sowa, John F. "Representing Knowledge Soup In Language and Logic" presented at the Conference on Knowledge & Logic, 15 June 2002.
32. Sowa, John. Ontology Web Site, http://users.bestweb.net/~sowa/ontology/
33. Watt, S. N. K. Seeing things as people: anthropomorphism and common-sense psychology, Unpublished PhD thesis, The Open University. (1997) available online at www.comp.rgu.ac.uk/staff/sw/publications.htm
34. Woods, D. D., Roth, E. M. & Bennett, K. B. Explorations in Joint Human-Machine Cognitive Systems. In W. Zachary & S. Robertson (Eds.), *Cognition, Computing and Cooperation.* Norwood, NJ: Ablex Publishing, 123-158, 1990
35. Zadeh, Lofti *In Quest of Performance Metrics for Intelligent* Systems – A Challenge that Cannot be Met with Existing Methods, PerMIS 2002.
36. Zadeh, Lofti Protoform Theory and Its Basic Role in Human Intelligence, Deduction, Definition, and Search PerMIS 2003.